White Paper: Future Bandwidth requirements for student housing networks

Andrew Marshall, FBCS, CITP, CNP
Campus Technologies Inc
July 2011

Campus Technologies Inc
4119 Walnut Street, Suite 200
Philadelphia PA 19104
www.campustechnologies.com
amarshall@campustechnologies.com
(215) 243-5100

Copyright

Revision

Issue	Date	Author	Comments
1.0a	12/14/2007	Andrew Marshall	Initial Draft
1.0b	12/15/2007	Andrew Marshall	Second Draft
1.0	12/17/2007	Andrew Marshall	Issued document
2.0a	8/2/2010	Andrew Marshall	Second Edition Draft 1
2.0b	7/7/2011	Andrew Marshall	Second Edition Draft 2
2.0	7/16/2011	Andrew Marshall	Issued document

Audience

This white paper is intended for use by Campus Technologies Inc and its partners and affiliates, and to selected and authorized third parties on a case-by-case basis.

Contents

Executive Overview

Internet bandwidth available to end-users of networks will grow over the next ten years, and the effect of this growth will be to render ineffective or obsolete many student housing networks in a relatively short time period. Many off-campus student housing communities will be disadvantaged by having limited or no ability to change their current networks due to inflexible term contract arrangements.

This is the second revision of this document, the original having been created in December of 2007. Actual data reported between the original document and today validates the original document, and we have updated this revised version with the most up to date information. We have also shifted focus from visitor based and subscriber networks to Student Housing networks as these are our specific area of interest, although the conclusions and data are largely transferable between the two areas.

The growth in Internet bandwidth available[1] to end-users varies dramatically by country. A September 2010 OECD study of thirty four countries fastest advertised speed shows that the difference between the highest (Sweden, 1,024 Megabits per second per user) and the lowest (Mexico, average 10 Megabits per second per user) is substantial. The US ranks 24th in the study with a maximum of 51 Megabits per second per user. This represents a US ranking decline from 14th in our original December 2007 paper. In a comparison of average download speeds, the US fares even worse at 30th out of 34.

Some countries (Sweden, Slovenia, Slovak Republic, Portugal) now have in production gigabit (1,024 megabits per second) services available to residential and commercial users and 100 megabits per second or greater services are available in twenty-three of the thirty four countries surveyed. 100 megabits per second consumer services were expected to be under limited trial in the United States in our original paper but that has not yet materialized[2].

Previously, a law termed Nielson’s law stated that bandwidth available to high-end-users will increase at a rate of 50% each year. While this has proved largely correct in the last decade in the US, some other countries have outstripped this law and are continuing to do so. In the case of Hong Kong we previously noted that instead of 50% the actual number in 2005 was 1,000%. We put forward for consideration that the next decade growth may outperform, but is unlikely to underperform Nielson’s law on a global basis, and that the use of Nielson’s law is the most conservative formula that should be used to state future bandwidth needs, and then it should only be used with an understanding of the risk of possible understatement. We would cite Portugal, Slovenia, the Slovak Republic and Sweden who have exceeded Neilson’s law by orders of magnitude in the last three years. In addition we warn that the change in media consumption habits is leaning away from traditional media such as cable and satellite TV and moving more towards real time streaming on the public Internet. This has the potential to substantially accelerate growth of demand for bandwidth.

[1] Download bandwidth only

[2] Campus Technologies Inc. introduced production 100 Megabit to the user services in mid 2010 but as this is restricted to student housing at present it has not been counted.

Using Nielson’s law in the interest of conservatism, we calculated the average user and high-end-user bandwidths for the US in the next ten years.

We analyzed these projections and concluded that student housing networks will be impacted by being unable to keep up with expected average user bandwidths within a very short period of time, in the case of gateway components in some networks as little as a year. In a specific example, a 500 user student housing network gateway performing concentration or oversubscription that currently only supports 100 Megabits per second of throughput will be unable to sustain the expected subscriber throughput in the very near future and as a result may seriously impact the service quality provided to end-users. In many, if not all, off-campus student housing communities, degradation in Internet service below that of competing properties places the community at a significant leasing disadvantage by reducing the popularity of the community with potential residents.

We further concluded that it is essential that network designers and operators should consider these bandwidth growth projections in the interest of investment protection. Specifically, the lifespan of 100 Megabit Ethernet switches and Category-5 or 5E wiring infrastructure should be viewed as limited, in the case of the US only surviving unaffected until sometime between 2011 and 2014, and with varying but identifiable end-of-life projections for other countries. Networks under planning or construction today should consider building to the Category 6 specification for wiring and ensure that distribution, core and gateway electronic components can manage the workload for their expected lifespan without limiting the bandwidth available to each user in their respective markets.

We further conclude that no networks should be designed, deployed or upgraded with 100 Megabit only Distribution or Core layers or Gateway components from now on.

The Bandwidth landscape and its future

Bandwidth available to end-users

This section refers to bandwidth available to end-users. To be specific, this is defined as the downstream bandwidth advertised as available to that sole individual user from that users' bandwidth provider for either commercial or residential use. Wherever bandwidth is referred to in this document, this definition is used unless explicitly stated otherwise. In student housing terms, think of it as bandwidth available 'to the bed'.

As of July 2011, there are a wide variety of available 'broadband' or 'high speed' options available to consumers, but these vary dramatically by country. In a study conducted by the OECD[3] in September 2010, the fastest advertised bandwidth available to consumers by country was plotted for 34 countries. The results are summarized in the graph below:

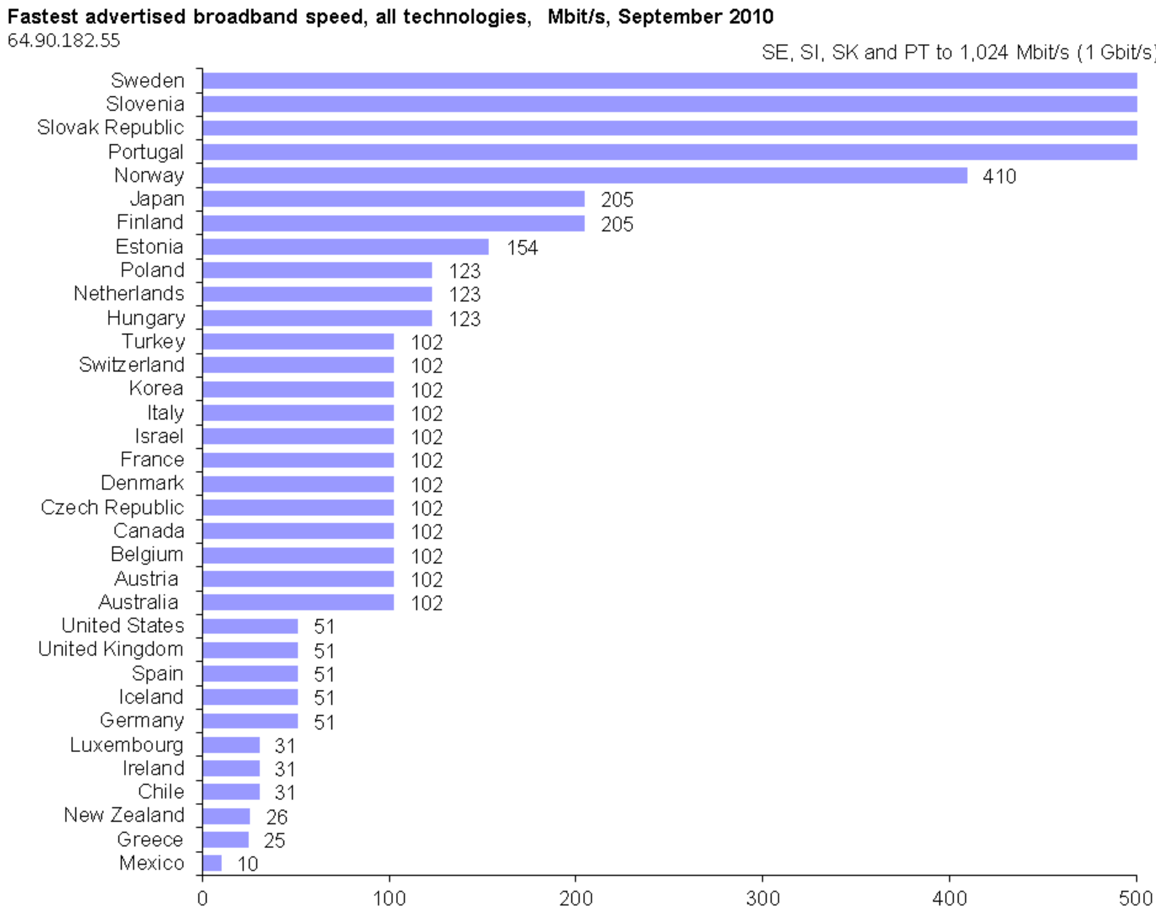

(source: OECD)

[3] Organization for Economic Co-operation and Development, www.oecd.org

As can be seen the state of the art in terms of national advertised bandwidth available to the end-user is in the 100 Megabit per second range or better, with the US at 51 Megabits. There are a number of factors that influence variances between countries, however this topic is outside the scope of this document. The intent of taking a snapshot with independent data is to show the following:

- The state of the art for maximum bandwidth available to end-users today;
- The scale of advances in this being made in some geographies, and
- To verify and validate predictions for future bandwidth use, which in turn will determine the requirements for delivery systems and network components.

For completeness, in a plot of <u>average</u> download bandwidth, the US fares worse still at 30th place at 14.6 Mbit/sec download speed.

Average advertised broadband download speed, by country, kbit/s, September 2010

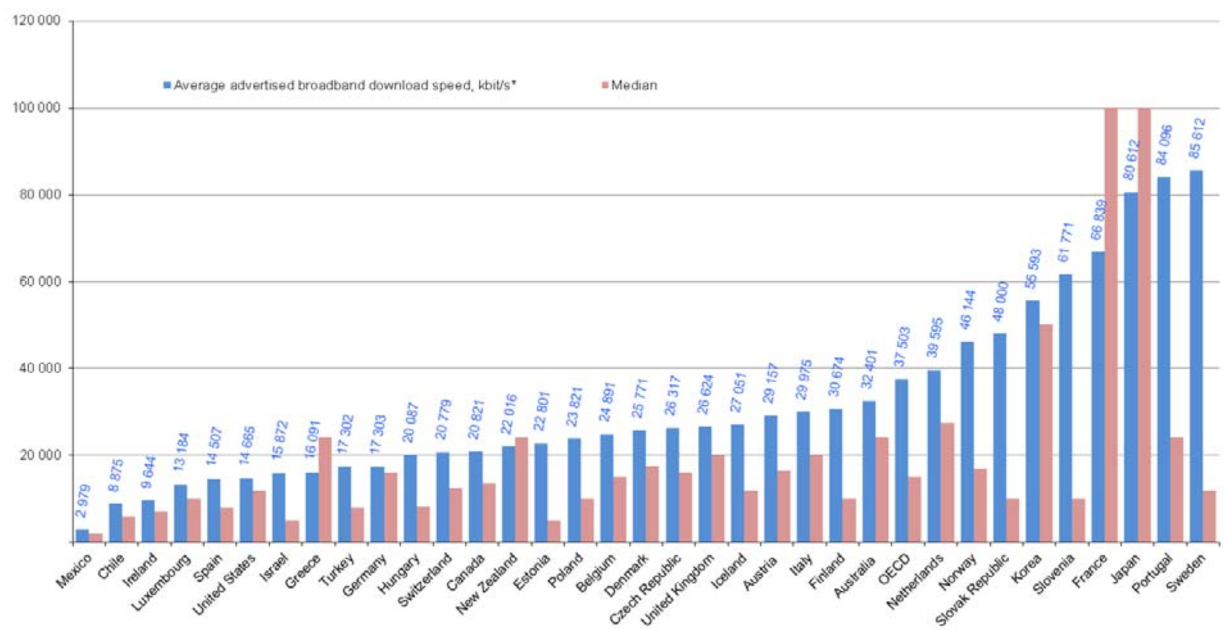

The global state of the art, or indeed the norm, is well in excess of current US Internet provisioning and that this demonstrates that the technology to deliver these bandwidths to end users is available and in common production use. This is a useful predictor of where the US is headed, and the numbers shown when compared to our original 2007 paper support our predictions.

Bandwidth available to end-users is increasing

There is incontrovertible evidence to demonstrate that over the last ten years that demand for greater bandwidth by subscribers in the US has been increasing, and that bandwidth available and delivery mechanisms to users to satisfy those demands has been increasing, although not necessarily at the same level. We would speculate that there is significant latent demand in the US for additional bandwidth and improved Internet performance.

Nielsen's law, from 1998[4], one of the only rules of thumb available for estimating Internet bandwidth demand, states that: "a high-end-user's connection speed grows by 50% per year"

The key drivers behind the spiraling demand for increased bandwidth by end-users are well documented and that demand trend is mainly self sustaining. The expectation of end-users is increasing commensurately and a wide variety of rich media applications are now de facto standards and part of everyday culture.

> In student housing in particular, the Gen-Y demographic has a high expectation that they will be able to consume Internet bandwidth on demand. Even in the last few years, the consumption of technology by student housing residents (and indeed by many other Internet users) has directly contributed to demand growth and supply challenges. Below are some examples:
>
> **Video and TV consumption**: This is shifting from traditional broadcast TV and rented physical videos significantly in the last five years. Today's student housing resident is consuming video on a variety of displays, from laptops to large screen TV's where the content has been sourced using an Internet connection. Many TV's and DVD's sold today have Ethernet jacks for connecting to the Internet, and specialized media players such as the Roku[5] depend on a reliable connection. As media moves from standard to high definition, the bandwidth requirement will increase dramatically. The nature of streaming video is such that multiple users in the same community will consume bandwidth each time one of them downloads or streams a video.
>
> **Music and audio**: This is almost exclusively streamed or downloaded using Internet radio and many music download sites. The recently announced initiatives that place users content 'in the cloud' shifts usage from one time download to constant streaming.
>
> **Video Gaming**: The trend here is shifting from one or several people using an Xbox or similar console in the same room to online gaming in real time with players anywhere in the world. This demands bandwidth, stability, and fast response times.
>
> **Other**: personal videos and photos are all stored online in services such as Flickr, much college research once confined to libraries is now in the domain of search engines. Social media is now an 'always-on' experience for most college students.

In a study conducted by Sandvine[6] in May 2011, statistics gathered show a major shift in the type of traffic being used on a residential Internet connection in the US. The largest overall proportion is Netflix (streaming video) with a massive 24.71% of utilization. Even more

[4] http://www.useit.com/alertbox/980405.html
[5] http://www.roku.com
[6] http://bit.ly/kWf447

noticeable is that 'traditional' traffic accounts for only 17.18% of overall bandwidth utilization.

	Upstream traffic		Downstream Traffic		Total Traffic	
Rank	**Application**	**Share**	**Application**	**Share**	**Application**	**Share**
1	BitTorrent	52.01%	Netflix	29.7%	Netflix	24.71%
2	HTTP	8.31%	HTTP	18.36%	BitTorrent	17.23%
3	Skype	3.81%	YouTube	11.04%	HTTP	17.18%
4	Netflix	3.59%	BitTorrent	10.37%	YouTube	9.85%
5	PPStream	2.92%	Flash Video	4.88%	Flash Video	3.62%

Examination of available technologies and available services over the past decade in the US shows that the law postulated by Nielsen in 1998 has historically proven correct and that the high end-user's bandwidth available has increased by approximately 50% each year.

Notwithstanding historical performance, based on all available data it would not be unreasonable to conclude that Nielsen's law is fast being rendered obsolete by the growth of previously unforeseen traffic such as video and that bandwidth available to end-users, especially in certain geographies, is far outstripping that rate that would be expected using the law. We would cite as examples the cases of HKBN, the Hong King Broadband Network, which has been providing 100 Megabits per second to the home since 2004 and 1,000 Megabits per second to the home since 2005, not an experimental network but a major deployment that passes 800,000 homes[7], and the recent deployment in several countries of gigabit and multi-hundred megabit services, exponential increases over the accepted current norms in the US. Other countries are deploying or planning to deploy 100 Megabit per second and 1 Gigabit per second bandwidth to the end user. These examples serve to reinforce our assertion that bandwidths available over the next ten years will continue to increase at least at the rate proposed by Nielsen's law, and probably considerably greater. The US is unlikely to be left very far behind despite the logistical challenges of infrastructure deployment in such a large country, and Verizon has conducted trials of a 100 Megabit per second service using its FiOS deployment[8] although this has yet to materialize in production. In a widely publicized release, Google have announced[9] a test deployment of gigabit residential networking in Kansas City to test the next generation of technologies, and this is expected to go live in 2012.

> *"The only way you're going to win is to have capacity that nobody else can beat," Verizon Communications Inc. Vice Chairman Lawrence Babbio*

The projections used elsewhere in this document apply the conservative approach of applying Nielsen's law to the 2010 average data researched by the OECD and our own network utilization measurements.

[7] http://www.hkbn.net/bb1000/
[8] http://fiberopticnews.blogspot.com/2007/10/fios-is-catching-on-in-dallas-and.html
[9] http://www.wired.com/epicenter/2011/03/google-fiber-kansas

Estimates of future per user bandwidth requirements for 2010 – 2020

Using Nielsen's law, we have extrapolated the domestic US available bandwidth per end-user based on the law's original precept: that the bandwidth available to a high end-user will increase by 50% per year.

Year	Rate	Scale
2010	**51**	**Mbps**
2011	77	Mbps
2012	115	Mbps
2013	173	Mbps
2014	259	Mbps
2015	388	Mbps
2016	582	Mbps
2017	873	Mbps
2018	1,310	Mbps
2019	1,153	Mbps
2020	1,965	Mbps

Table 1 – High end US User extrapolation

In this extrapolation, 100 Megabit per second per user demands are exceeded by 2012, and Gigabit per user demands by 2018. We should point out many other countries are already at the point that the USA is predicted to reach by 2012, and that some are already at the point that the USA is expected to reach by 2017/18.

In the next table, the *average* US speed quoted in the 2010 OECD report of approximately 14 Megabits per second per user, the following extrapolation results:

Year	Rate	Scale
2010	**14**	**Mbps**
2011	21	Mbps
2012	31	Mbps
2013	47	Mbps
2014	71	Mbps
2015	106	Mbps
2016	159	Mbps
2017	239	Mbps
2018	359	Mbps
2019	538	Mbps
2020	807	Mbps

Table 2 –Average US User extrapolation

In this table 100 Megabit per second per user demand is reached in 2015.

We should like to point out once again that although these tables demonstrate future dates for attaining the 100 Megabit per second per user demand barrier that many countries have already attained this point, and that several other countries will meet or exceed the 100 Megabit per user point prior to 2012/3. We would argue that this is an indication that

Nielsen's law may be becoming obsolescent, however for the purposes of this document and in the interests of conservatism we will use the Nielsen method. We would however caution that we believe these numbers to be possibly understated for planning purposes.

Student Housing networks perform concentration

In most student housing networks, an amount of bandwidth is distributed among a number of end-users. Similar networks can include, but are not limited to Enterprises, Multi Family residential units, Multi Tenant commercial units, Educational campuses and student housing, Hospitals, Hotels, Military housing, Conventions and meeting centers, Hotspots, transportation hubs and vehicles and many more.

Typically in this kind of network an amount of external bandwidth is connected to a gateway device that manages bandwidth shaping and provisioning of internal end-users in these networks. It is normal industry practice to apply a concentration ratio or oversubscription rate; i.e. to provision 10 users with a 10 Megabit per second per user demand rate you would not provision a gateway external connection that is simply 10 x 10 = 100 but instead use the formula (Number of users x required per user bandwidth / concentration factor). Thus a concentration factor of 10 (10:1 concentration) would require an external bandwidth connection of 10 Megabits per second in our previous example. These gateway devices usually also provide rate limiting or bandwidth shaping functionality to ensure fair distribution of resources between active users.

Concentration is made possible by the fact that historically users rarely if ever simultaneously demand the full bandwidth available to them and that traffic patterns such as web browsing are inherently asynchronous with pause times between page demands. Clearly however, the concentration ratio must be chosen carefully as too high a ratio will result in user slowdown as users are unable to obtain adequate bandwidth, and too low a ratio will result in idle bandwidth and excessive costs. Obtainable concentration ratios can vary substantially by class of user, however for this example we are using a nominal 30:1[10].

In this manner a network with 200 users being allocated the US average of approximately 14 Megabits per second per user will require external bandwidth of 93 Megabits per second on their external connection, i.e. the connection to the community.

Using this data, and applying the growth rates specified in the possibly understated Nielsen's law, the throughput that must be handled by the network over the next ten years is as shown in table 4. Note the baseline date in use is 2010 as the OECD data used originates from September 2010.

---

[10] We must caution that the advent of more streaming requirements will reduce the concentration rates achievable as time progresses and the traffic becomes less asynchronous. We previously noted that traditional traffic is already a minority and this will significantly impact achievable concentration ratios and oversubscription in the future as this trend develops.

Year	200 users	500 users	Scale
2010	**93**	**233**	**Mbps**
2011	140	350	Mbps
2012	210	525	Mbps
2013	315	788	Mbps
2014	473	1,181	Mbps
2015	709	1,772	Mbps
2016	1,063	2,658	Mbps
2017	1,595	3,987	Mbps
2018	2,392	5,980	Mbps
2019	3,588	8,970	Mbps
2020	5,382	13,455	Mbps

Table 4 – Gateway and network bandwidth extrapolation

Effects of bandwidth futures on access networks

In the preceding sections we demonstrated that the growth of bandwidth available to the end-user on a per-user basis, when determined on a highly conservative basis, will exceed the capacity of today's student housing, subscriber and visitor based networks in a relatively short period of time. Property owners and managers and their network designers and planners need to take into account the demands on, and potential lifespan of various components of their wired networks, and also need to have visibility of changes that will likely be required within a relatively short timeframe.

From the calculations in table 4, it is clear that networks of 500 users or above that can only deploy 100 Megabit Ethernet technology at the Gateway (or deploy Gateway devices that can only sustain loads of 100 Megabits or less) and/or at the core will need to either reduce service levels to users by restricting bandwidth or increasing concentration levels, or segment the network to a more granular level and deploy multiple gateways. This will only provide temporary relief as the continuing increase in per user bandwidth expectation will overrun any such temporary measures in a short time horizon.

In this section, we discuss the implications of the conclusions of the previous sections on the physical plant and wiring, the access layer, the distribution/core layer, and access gateways.

Physical plant and wiring infrastructure

Most new personal computers sold today contain Gigabit capable network interface cards, and these cards are capable of also operating at 100 Megabits per second by auto sensing the capabilities of their connection. Given the usual lifetime of a computer of 2-4 years, and the fact that most of these devices can easily be upgraded to Gigabit Ethernet, it is a reasonable assumption that the majority of devices connection to networks within the next four or five years will be Gigabit capable.

Most non-Enterprise Ethernet access networks will continue to operate at 100 Megabits per second in the short term and the standardization of Network Interface Cards on Gigabit Ethernet will not affect this. The challenge lies in determining the life expectancy of the 100 Megabit per second capable Category-5E standard[11] for structured wiring. In the USA, this will become a sub-standard capability between 2011 and 2013 according to our findings. The international lifespan of Category-5E will be less, on average (see table 3), however this indicates that half of all countries outside the US will overrun Category-5E prior to that date.

Conservative design would indicate than any new networks being built today should consider their life expectancy for wiring infrastructure within these parameters and utilize Category 6 structured wiring[12] to avoid costly wiring upgrades in the short to medium term. In addition, operators and owners of current network should begin determining when they will need to upgrade their physical infrastructure and plan accordingly.

---

[11] ANSI/TIA/EIA-568-B.1-2001, -B.2-2001, and -B.3-2001
[12] Also defined in ANSI/TIA/EIA-568-B.1-2001, -B.2-2001, and -B.3-2001

Wireless Access

Wireless Access points and wireless support for mobile computing, while being outside the direct scope of this paper present unique challenges. Most, but not all local area wireless access technologies such as those defined in the 802.11x standards rely on connection to access points that require wired infrastructure and access or distribution layer capability in the same way as wired connections and similar challenges will apply. Wireless connections continue to provide slower connections than wired connections, and we believe that trend will continue for some time to come.

Wide area wireless access technologies, while important in their own right, are likely to remain in the mobile computing domain for the time being and are not considered to have material bearing on the topic of this paper.

The Access Layer, and Access Layer switches

Using the definitions of the Cisco hierarchical three layer switching model[13], the Access Layer connects individual users to the delivery network. Currently, many subscriber and visitor based networks use Ethernet switches in the access layer with a maximum capability of 100 Megabits per second per port. In the USA, for the time being, we believe that this remains a viable option, as the lifespan of the electronics in the access layer is usually less than that of the wiring infrastructure and will be adequate for the next three years based on the data in Table 2. We would make two provisos: firstly that the connection between the Access Layer switch and the Distribution or Core Layers should be Gigabit speed as a minimum to allow aggregation of traffic in line with the demands outlined in this paper, and secondly that the network is intended to serve average rather than high end-users. If the network is intended to serve high end-users the lifespan of 100 Megabit ports will be less as shown on Table 1.

In non-US deployments the determination of the suitability of 100 Megabit per second per port will be a function of local bandwidth growth and would need to be determined on a case by case basis. Broadly, the countries surveyed by the OECD with an average advertised bandwidth higher than the USA will get less operating life from a switch at this speed and should consider Gigabit per second per port switching infrastructure sooner in the Access Layer. Those less than the US average may obtain longer service life from 100 Megabit per second commensurately.

The Distribution and Core Layers and associated switches

Once again using the definitions of the Cisco hierarchical three layer switching model, the Distribution and Core Layers may be separate or combined depending on the individual network design.

Any port in the Distribution/Core layer that connects an Access Layer switch or switches must have the ability to handle the traffic passed to and from those switches. Based on the

[13] http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2002.htm

bandwidth demands outlined in this paper, in the majority of cases networks being designed and deployed today should utilize Gigabit ports as a minimum and have an adequate aggregate throughput capability.

Gateway and Core throughput requirements

In a student housing, subscriber or visitor based network, access to the network by the user is usually managed by a gateway device that provides configuration isolation, bandwidth shaping, access control, billing support, proxy, firewall and DHCP/NAT functionality.

The total network throughput will determine the sizing of the Core Layer switching and Gateway capabilities.

Using this data, and applying the growth rates specified in the somewhat conservative Nielsen's law, the throughput that must be handled by the core and gateway components network over the next ten years are as in table 4, repeated as follows:

Year	200 users	500 users	Scale
2010	**93**	**233**	**Mbps**
2011	140	350	Mbps
2012	210	525	Mbps
2013	315	788	Mbps
2014	473	1,181	Mbps
2015	709	1,772	Mbps
2016	1,063	2,658	Mbps
2017	1,595	3,987	Mbps
2018	2,392	5,980	Mbps
2019	3,588	8,970	Mbps
2020	5,382	13,455	Mbps

Table 4 – Gateway and network bandwidth extrapolation

From these calculations, it is clear that networks of 500 users or above that only deploy 100 Megabit Ethernet technology at the Gateway and the core will need to either reduce service levels to users by restricting bandwidth, allowing contention and thus increased latency, increasing concentration levels, or segmenting the network to a more granular level and deploy multiple gateways. These tactics will, however, only provide temporary relief as the continuing increase in per user bandwidth expectation will overrun any such temporary measures in a short time frame.

Even in 200 user networks the life expectancy of existing gateways with 100 Megabit per second throughput is less than a year. In larger networks, saturation exists already and users attempting to employ 100 Megabit per second core and gateway technology will be unable to provide national average performance per user due to exceeding their specifications.

Other factors

Technologies

We believe that the rapidly growing trend of accessing streaming multimedia using the public Internet will give rise to new bandwidth conservation technologies. As a background, media is streamed in a unidirectional manner (towards the end user) and in what is termed ‘unicast’ mode. This means that if four people watch the same video, the bandwidth required is four times the requirement of one person as each independent stream requires Internet bandwidth. In this way, Internet streaming is much less efficient use of bandwidth than traditional analog or digital TV signals that broadcast a single signal to multiple recipients. Although internet technologies exist to support multicast viewing, they are not commonly adopted or deployed.

This is a dilemma: on one hand consumers are demanding the convenience of Internet based streaming, on the other it is a highly inefficient use of resources.

We predict that large multi user networks such as student housing networks will need to deploy bandwidth conservation techniques that keep multiple demands for identical content from consuming all available external Internet bandwidth. Such techniques are outside the scope of this document; however we have several deployments testing bandwidth conservation technologies in our own networks and are convinced that this will be a key factor to succeeding in facing growth challenges.

Contractual challenges

Most student housing communities are locked into term contracts for long periods, with limited control over the quantity of bandwidth currently being provided to those communities, and limited or no ability get out of those contracts. This is the single largest risk factor facing student housing owners and operators today; a lack of ability to control the quality and quantity of a critical service that will directly affect leasing and NOI. We would recommend that all owners and operators carefully examine the contract position of their assets with regard to Internet provisioning and maintain a high level of flexibility in future years.